

Statistical Signal Processing with Physics-Based Models: UXO Detection and Identification: CU-1123

Performing Organization: Duke University
Principal Investigator: Leslie Collins

Final Progress Report
December 23, 2002

I. Project Background

Using current technologies, the cost of identifying and disposing of UXO in the United States is estimated to range up to \$500 billion. Site specific costs range from \$400/acre for surface UXO to \$1.4 million/acre for subsurface UXO. There are 1900 Formerly Used Defense Sites (FUDS) and 130 Base Realignment and Closure (BRAC) installations that need to be cleared. Several sensor modalities are currently being explored for the detection and identification of surface and buried UXO. These include electromagnetic induction (EMI), magnetometers, radar, and seismic sensors. These sensors experience little difficulty detecting the UXO, thus detection does not create the bottleneck that results in the high cost of remediating sites. The primary contributor to the costs and time associated with remediating a UXO contaminated site is the high false-alarm rate associated with each of the sensors when operated individually.

II. Objective

In this project, we investigated the phenomenological aspects of the UXO detection, location, and discrimination problem using EMI, radar, and magnetometer sensors. The fundamental insight garnered by characterizing the underlying physics was transitioned into performance bounds, as well as high-performance sensor fusion and signal-processing algorithms for enhanced detection, location, and discrimination of buried UXO under a wide range of environmental conditions. The signal processing algorithms that were developed were evaluated on data collected in the field during two government-sponsored sensor demonstrations. Finally, the algorithms were applied to a “blind” data set and performance was assessed by an independent agency.

III. Technical approach

The technical approach that was pursued employed synergistic research activities in modeling, signal processing, and sensor fusion. We performed phenomenological modeling of wave propagation and scattering for ultra-wideband (UWB) radar and EMI sensors. Our phenomenological studies were performed in collaboration with SERDP-supported sensor-development programs underway in these areas (at NRL, ARL, and AETC). Our previously developed models were extended to allow arbitrary numbers of soil layers, arbitrary target shape and orientation, and to accurately account for all interactions. Using these models, we quantified the target types, depths, and soil conditions for which radar is an appropriate sensor. We also quantified the bandwidths

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 23 DEC 2002	2. REPORT TYPE		3. DATES COVERED 00-00-2002 to 00-00-2002		
4. TITLE AND SUBTITLE Statistical Signal Processing with Physics-Based Models: UXO Detection and Identification			5a. CONTRACT NUMBER		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Duke University, 2138 Campus Drive, Durham, NC, 27708			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 22	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

required to excite resonances for EMI and seismic sensors. In addition, we developed simpler phenomenological models for EMI and magnetometer sensors that could be used within the context of data inversion for object parameters.

These models of the wave physics, coupled with models of target, clutter, and environmental uncertainties, were incorporated into a statistical signal processing framework, thus novel, state-of-the-art optimal detection and identification algorithms were developed for each sensor. Bayesian algorithms, which provide the optimal solution to detection and identification problems, were investigated along with an algorithm based on a Hidden Markov Model formulation which is specifically suited for classification using data from multiple aspect angles. Algorithms based on Support Vector Machines (SVMs) were also considered to mitigate the necessity of estimating probability density functions. This estimation is an inherent step in Bayesian approaches. Finally, we developed sensor-fusion techniques that simultaneously exploited the richness and diversity of the phenomenology underlying multiple sensor modalities. In all cases, the algorithms that were developed were tested on data collected using sensor systems also under SERDP support, such as NRL's MTADS system, the ARL Boom-SAR, and Geophex's GEM-3 EMI sensor.

IV. Summary

For both EMI and SAR sensors, the wave-based phenomenological models that were developed provided predictions that matched measured sensor data extremely well. In the case of the SAR data, these predictions were used within a template matching construct (or correlation receiver) to improve the false alarm rate over RCS-based signal processing approaches. The EMI models were utilized to validate our more simplistic phenomenological models and to help demonstrate the rationale behind failures in the more simplistic models. For example, for complicated UXO, it was demonstrated that a simple single-dipole model could not be utilized to accurately predict sensor data from all possible target/sensor orientations.

Single-sensor and sensor-fusion algorithms were developed for EMI and magnetometer data based on inversion of measured data using the simple phenomenological data. The parameters estimated during the inversion process were submitted to several classes of algorithms including Bayesian, clustering, neural networks, and support vector machines. The Bayesian and SVM algorithms provided the most improvement in false alarm rates and were consistently the most robust. Data from both the JPG-IV and JPG-V demonstrations was processed. In all cases the algorithms developed under SERDP support out-performed those of the demonstrators. Based on the discrimination performance obtained on the JPG-IV data and the blind test results obtained on the JPG-V data, the statistical algorithms developed in conjunction with this project provided up to a factor of 3 improvement in false alarm rates.

V. Project accomplishments

Project accomplishments are divided into sections on modeling, performance bounds, and signal processing. Details regarding accomplishments are provided in the manuscripts included in this report in Appendix A.

Modeling

With regard to SAR-based systems, we have developed a Method of Moments (MoM) as well as a multi-level fast-multipole algorithm (MLFMA) for modeling radar scattering from electrically large targets arbitrarily buried in a half space. Although the MoM approach can handle large targets, such as UXO, in principle, the memory requirements and computation time become excessive. There has consequently been significant interest recently in the development of a new generation of fast algorithms to cover the range of frequencies for the problem of interest. In general, high frequency asymptotic techniques, e.g. physical optics (PO) are either not applicable or difficult to implement. The MLFMA model implemented under SERDP support results in considerable computational savings. A comparison between the MoM, PO, and MLFMA predicted response for a 155mm shell buried 2.54 cm beneath dispersive Yuma soil with a depression angle of 30 degrees is shown in Figure 1. Clearly, the response predicted by the MLFMA approach is consistent with that of the rigorous MoM scattering response. Figure 2 shows a comparison between the CPU and RAM required to calculate the response of each of the models as a function of number of unknowns, which is proportional to the size of the object. As the size of the object increases, the benefit of using the MLFMA model becomes increasingly apparent. The agreement with data measured by the Army Research Laboratory is also excellent. An example of the comparison between the MoM predicted radar cross section and that measured by ARL for the same 155 mm target is shown in Figure 3 located on top of Eglin soil.

The initial MLFMA implementation was applicable to perfectly conducting targets, and was limited in that the target could not touch the air-soil interface. We have extended the MLFMA model to the case of general dielectric targets in the presence of a half space, and the model has been modified such that the target can touch the interface. This allowed us to model a variety of clutter objects, including vegetation, trees, rocks, etc. and to begin to develop statistical clutter models for use with the radar signal processing. Figure 4 plots electric currents induced on a tree model, situated above a lossy half space (soil). These results were computed via Duke's multi-level fast-multipole algorithm (MLFMA). The tree is 6 m tall, and the results are for 500 MHz operation and vertical polarization. The currents are normalized with respect to the incident fields. Figure 5 plots the computed SAR imagery from the tree trunk considered in Fig. 4. The imagery is shown for the center of the SAR aperture position at five different azimuthal positions with respect to (the same) tree model. The results are shown for VV polarization, for a 50-500 MHz bandwidth SAR. These results underscore the strong aspect-dependent signature associated with vegetation. This aspect dependency in the signature is utilized within the context of the HMM-based signal processing.

The simpler EMI models that were developed can be applied to either time- or frequency-domain EMI data. We assume that at EMI frequencies, and for distances relatively far from the target, the induced magnetic fields can be well characterized by a dipole model. We first introduce a magnetic polarizability matrix that is unique to each target, and note that the eigen coordinate system is related to target geometry. We have shown that the response measured by a receiver can be calculated easily in terms of the canonical target response and weighting matrices that describe the target/sensor orientation and the object orientation. This model is similar to the dipole model proposed by AETC, but differs in that it specifically relates the polarizability matrix expressed in eigen coordinates to the characteristic decaying exponential functions (poles/modes) associated with a target. These modes were also calculated using the Method-of-Moments (MoM) code that was originally developed under the MURI, but was extended under SERDP support to consider arbitrarily large objects that are not required to be a body of rotation.

Under MURI support, we had developed a rigorous method of moments (MoM) model that allows computation of such modal poles. While accurate, the model was computationally intensive and yielded little physical insight beyond the poles themselves. It was also limited as to the type of object (body of revolution) and size of object which could be computed. We completed a Finite Element Method (FEM) model that can be used to compute the resonances of general conducting objects, including ferrous objects. This model, while computationally complex, can be utilized to explore the phenomenology associated with EMI responses for arbitrary targets at arbitrary orientations. Figure 6 shows a comparison of the FEM modeling prediction with measured data. The bottom portion illustrates the excellent agreement between the measured and modeled EMI response as a function of frequency to a ferrous cylinder. The top curve plots the imaginary part of the response, and the bottom curve plots the real part of the response. The line plots the measured data and the symbols plot the model predictions.

Performance Bounds

During the first year, we derived the CRLB for estimating time-domain EMI features, or decay rates. In the second year, we extended that analysis to frequency-domain EMI sensors, and derived the CRLB for the elements of the polarizability tensor. Also in the second year, we developed an improved estimation procedure wherein the “poles” were estimated, and preliminary results indicated that this estimation procedure came closer to achieving the CRLB. We also studied the behavior of this estimator in great detail. The method improves the estimates by altering the shape of the objective function, and thus is not dependent on the particular optimization technique implemented. Across a broad range of decay rates, SNRs, and ratios of decay rates, the RMS error and bias of the decay rate estimates are consistently better when the signal is parameterized in terms of poles than in terms of decay rates. In addition, the results indicate a dependence on the ratio of the decay rates in the signal. As the ratio approaches 1, the exponential signals become more similar and thus more difficult to separate.

As an example, consider a signal that is the sum of two decaying exponentials, the upper of which is 20,000. A series of lower decay rates were chosen so the ratio of the two decay rates was 0.1, 0.3, or 0.7. The RMS error and bias of the decay rate estimates are shown in Figure 7 as a function of the noise variance, corresponding to total SNRs from approximately 35 dB to 135 dB. Fielded EMI sensors typically operate at an SNR within this range. For each combination of decay rates, the RMS error is shown in the left panel and the bias is shown in the right panel. The dotted lines indicate the results obtained using the decay rate signal parameterization and the solid lines indicate the results obtained using the pole signal representation. Parameterizing the signal in terms of poles consistently provides better decay rate estimates as measured both by the RMS error and the bias.

Signal Processing

Our signal processing efforts focused on HMM processing for SAR data, Generalized Likelihood Ratio Test and Support Vector Machine processing for magnetometer and EMI data, and sensor fusion. Magnetometer and EMI data collected by MTADSNRL, Geophex, and Naeva were all considered. Specifically, we focused on two data collections that were performed in conjunction with two Jefferson Proving Ground demonstrations: JPG IV and JPG V.

BoomSAR Processing. We used the MLFMA model discussed above for modeling radar scattering to develop a hidden Markov model for processing SAR data. In particular, the validated model allowed us to consider radar scattering from an arbitrarily oriented buried UXO. The SAR image varies as a function of the target-sensor orientation. We therefore partitioned the physics into a set of states, each state characteristic of a set of target-sensor orientations over which the associated SAR image is relatively stationary. We initially used a physics-based matching-pursuits algorithm to perform feature parsing, and although this worked fairly well, it could be confused by clutter. We therefore investigated the use of the aspect-dependent *computed* SAR images from our electromagnetic model to design Wiener filters. The algorithm performed extremely well for angle-entry targets and moderately well for surface and flush buried targets. The performance degrades, however, when the aperture of the Boom-SAR does not illuminate the broadside of a target, however the performance of the HMM was always better than any of the individual prescreeners. If the SAR was flown in a circle, these problems would be mitigated.

Figure 8 shows receiver operating characteristic (ROC) curves for two detectors, as applied to ARL SAR data measured at Yuma Proving Ground. The detectors were applied to regions specified via the prescreeners. Results are shown for an HMM classifier, and a simple correlation filter. The HMM, as discussed above, explicitly exploits the multi-aspect information provided by the SAR. In these results the HMM detector is designed to find surface and shallow-buried UXO (not deeply buried targets), motivated by the goal of UXO-range detection.

JPG IV Processing. Performance at JPG-IV was summarized in a variety of ways, but we compared the performance obtained using the False Positives (probability of stating “UXO” when non-UXO is present) and True Positives (probability of stating “UXO” when UXO is present) measures. We applied both Bayesian techniques and fuzzy clustering techniques to the data collected by Geophex with the GEM-3, NRL with the MTADS system, and NAEVA with the PROTEM system. Fuzzy clustering techniques were considered in addition to the Bayesian approach as the amount of training data was minimal, and the phenomenological models were not yet complete. Algorithms were trained in a “round-robin” fashion: the data set was partitioned into a training set and a testing set, performance was evaluated, then the data set was re-partitioned. In addition, a second training set was supplied by NRL and the algorithms were trained on this second set and tested on the JPG IV data.

The importance of utilizing physics-based models for feature extraction is illustrated in Figure 9. These results show receiver operating characteristic (ROC) curves obtained for the GEM-3 sensor during the JPG-IV demonstration. Probability of detection is plotted as a function of probability of false alarm. The green symbol shows the performance obtained by Geophex during the demonstration. The red curve shows the performance of a statistical algorithm operating on features developed in an ad-hoc manner, whereas the blue curve shows the performance of a statistical algorithm operating on features obtained by inverting the sensor data for features associated with a phenomenological (dipole) EMI model. Clearly, performance is substantially better utilizing either statistical algorithm, but is improved most when physics-based features are utilized.

We also investigated algorithms for processing both the *MTADS* EM and magnetometer data, as well as algorithms for fusing this data. We worked with Dr. Sean Hart and others at NRL to compare our algorithms to the Probabilistic Neural Network algorithm that has been developed there. Several signal processing algorithms were considered. First, we considered a statistical signal processing approach based on the generalized likelihood ratio test. Traditionally, such statistical processors are less robust when limited training data are available, or when the statistics of the parameters that affect the signal cannot be adequately characterized, as was hypothesized to be the case for the JPG IV demonstration. Therefore, sub-optimal but potentially more robust approaches were considered. These included a maximum likelihood estimation-based clustering algorithm whose clusters are initially determined using the ISODATA algorithm, and a subtractive fuzzy clustering technique.

The ROC curves for the algorithms described above are shown in Figure 10 for case where leave-one-out training was applied in conjunction with the JPG IV data. All processors utilize both magnetometer and EMI data, thus the results are essentially “sensor fusion” results. ROC obtained for a simple threshold processor operating on the EMI data alone. The EMI threshold is indicative of the performance obtained by a cueing processor that might be used as a pre-processor. In addition, the performance of the *MTADS* system using “man in the loop” processing as scored for the JPG IV demonstration is shown. Clearly, a simple threshold on the EMI data is not an effective discriminator of UXO from non-UXO items. The maximum-likelihood clustering

algorithm improves performance somewhat, and its performance is near that achieved by the NRL “man-in-the-loop” approach. The PNN, subtractive fuzzy clustering technique, and the Bayesian approach provide substantial performance improvements over that obtained by the maximum likelihood clustering technique. This may be a result of the additional information incorporated by the fuzzy rule base and the PNN, or the correlation structure incorporated into the Bayesian processor. Although the limited amount of data does not allow strict statistical comparisons across algorithms, it appears that the PNN performance is better than the other algorithms at the lower true positive rates and that the GLRT performance is better than the other algorithms at the higher true positive rates. The subtractive fuzzy technique tends to perform at the average of the PNN and GLRT algorithms. Similar, though slightly lower, performance trends were observed when a separate training set was utilized.

JPG V Processing. The JPG V test was more consistent with a traditional field survey in that data was gathered continuously instead of in a grid surrounding pre-defined locations where objects of interest were emplaced. Three areas were surveyed in JPG V and were designated Areas 1, 2, and 3. Ground truth was released for area 3, but was not released for area 2. Both the Bayesian approach, or Generalized Likelihood Ratio Test (GLRT) and the SVM were applied to data collected during the JPG-V demonstration by the GEM-3 and MTADS sensors respectively. Algorithms were developed based on the Area 3 data and then applied to Areas 1 and 2. The output of the algorithms for each anomaly defined by the vendor was provided to IDA for blind scoring. Performance is reported as probability of detection as a function of false alarm count, and all results are scored without the 20 mm ordnance items.

ROCs for the performance results are shown in Figures 11-13. The green star represents the performance obtained by the mag and flag crew, the blue line represents the off-site performance obtained by the vendor, and the black and red curves show the performance of the SVM and GLRT algorithms respectively. Similar to the results for the JPG-IV demonstration, the performance obtained with the statistical algorithms is substantially better than that obtained with more traditional approaches. In addition, the fact that performance gains were obtained under blind test conditions suggests that statistical processing approaches based on features extracted using phenomenological models can be used to decrease the false alarm rate.

VI. Conclusion

The hypotheses underlying this research were that phenomenological models could be utilized to facilitate statistical signal processing, and that statistical signal processing algorithms could be used to reduce the false alarm rates associated with discriminating UXO from anthropic and geological clutter. Both of these hypotheses were validated during the term of this research project. Statistical algorithms that were based on features extracted from EMI and magnetometer data using phenomenological models outperformed similar algorithms operating on ad-hoc features. Similarly, statistical algorithms out-performed more traditional algorithms operating on the same data sets and

using similar training sets. In particular, these hypotheses were validated in recent performance results from two different sensors in a blind field test. These results suggest that the algorithms developed under this program should be transitioned to the end-user to aid in reducing false alarm rates in realistic field scenarios.

VII. Transition Plan

In collaboration with AETC, we have obtained ESTCP support to roll the model, inversion, and processing algorithms into the GeoSoft platform. The transitioned algorithms will be tested on field data for verification purposes. Issues involving updating training sets and data validation will also be considered.

VIII. Recommendations

The ESTCP transition path is appropriate for moving the algorithms developed under this program to the end user. However, it is important to note that these algorithms were developed assuming a reasonable training set is available, prior knowledge regarding the targets present in a data set is available, precisely positioned data is available, and that the individual anomalies occur in isolation. Many of these assumptions are not always valid, and care must be taken to only apply the algorithms developed here when the algorithms are valid. Additional research must be pursued, and in fact in several cases is ongoing, in order to develop algorithms that are robust to these assumptions. In addition, as more information becomes available as to the variability of UXO signatures, this information should be incorporated into the formulation of the statistical algorithms.

IX. Appendix A

List of manuscripts supported by SERDP (attached). Technical abstracts follow manuscripts.

- (1) Rigorous Modeling of Ultrawideband VHF Scattering from Tree Trunks Over Flat and Sloped Terrain, He, Geng, Nguyen, and Carin, IEEE Trans. Geosc. Remote Sensing.
- (2) Multi-aspect Detection of Surface and Shallow-Buried Unexploded Ordnance via Ultra-Wideband Synthetic Aperture Radar, Dong, Runkle, Carin, Damarla, Sullivan, Ressler, and Sichina, IEEE Trans. Geosc. Remote Sensing.
- (3) On the Wideband EMI Response of a Rotationally Symmetric Permeable and Conducting Target, Carin, Yu, Dalichaouch, Perry, Czipott, Baum, IEEE Trans. Geosc. Remote Sensing.
- (4) Multilevel Fast-Multipole Algorithm for Scattering from Conducting Targets Above or Embedded in a Lossy Half Space, Geng, Sullivan, Carin, IEEE Trans. Geosc. Remote Sensing.
- (5) A Comparison of Algorithms For Subsurface Target Detection and Identification Using Time-Domain Electromagnetic Induction Data, Tantom and Collins, IEEE Trans. Geosc. Remote Sensing.
- (6) Performance Bounds and a Parameter Transformation for Decay Rate Estimation, Tantom and Collins, IEEE Trans. Geosc. Remote Sensing
- (7) A Comparison of the Performance of Statistical and Fuzzy Algorithms for Unexploded Ordnance Detection, Collins, Zhang, Li, Want, Carin, Hart, Rose-Pehrsson, Nelson, McDonald, IEEE Trans. Fuzzy Systems.
- (8) Sensing of Unexploded Ordnance with Magnetometer and Induction Data: Theory and Signal Processing, Zhang, Collins, Yu, Baum, and Carin, IEEE Trans. Geosc. Remote Sensing.
- (9) Unexploded Ordnance Detection Using Bayesian Physics-Based Data Fusion, Zhang, Collins, Carin, Integrated Computer-Aided Engineering.

X. Appendix B - Figures




Figure 1. Comparison between MoM and MLFMA predictions




Figure 2. Comparison of CPU and RAM savings of the MLFMA algorithm compared to the MoM approach. Calculations are for a SGI Origin 2000, 300MHz processor with aR12000 CPU chip.




Figure 3. Measured RCS Boom-SAR data (top) and model predictions (bottom)




Figure 4. Electric currents generated by the SAR on a modeled tree.

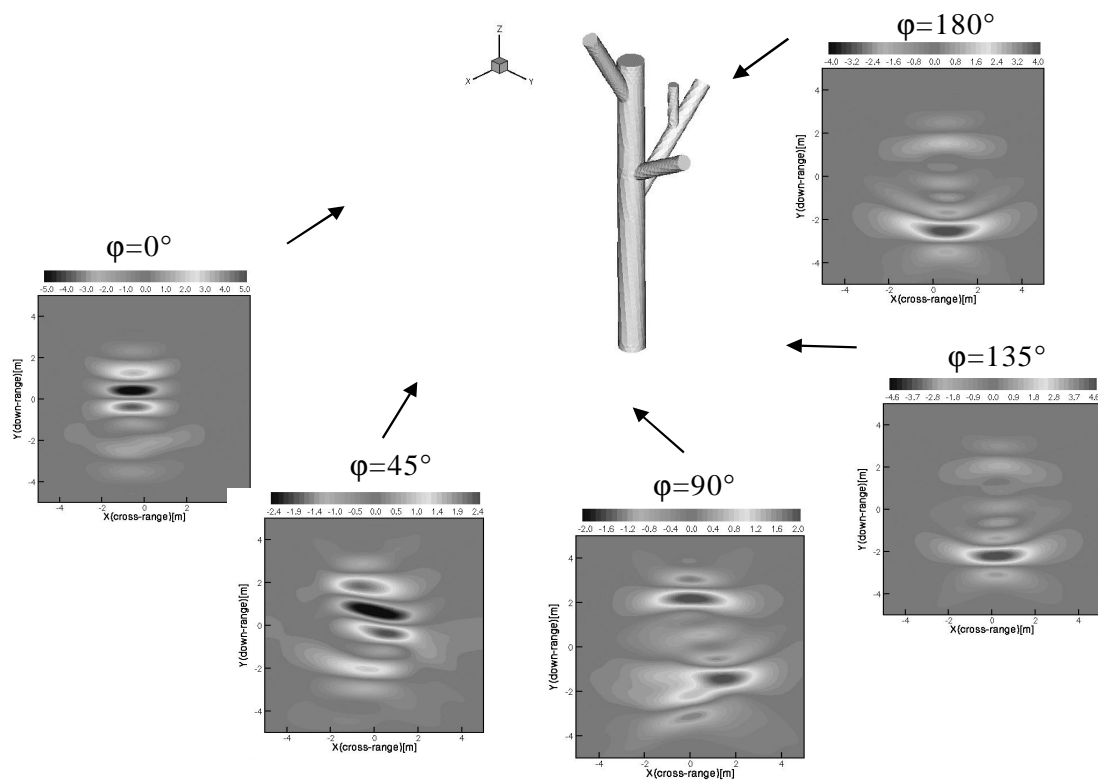


Figure 5. Modeled SAR imagery from the tree modeled in Figure 2 at various aspect angles.




Figure 6. Comparison of measured data and FEM EMI model predictions.

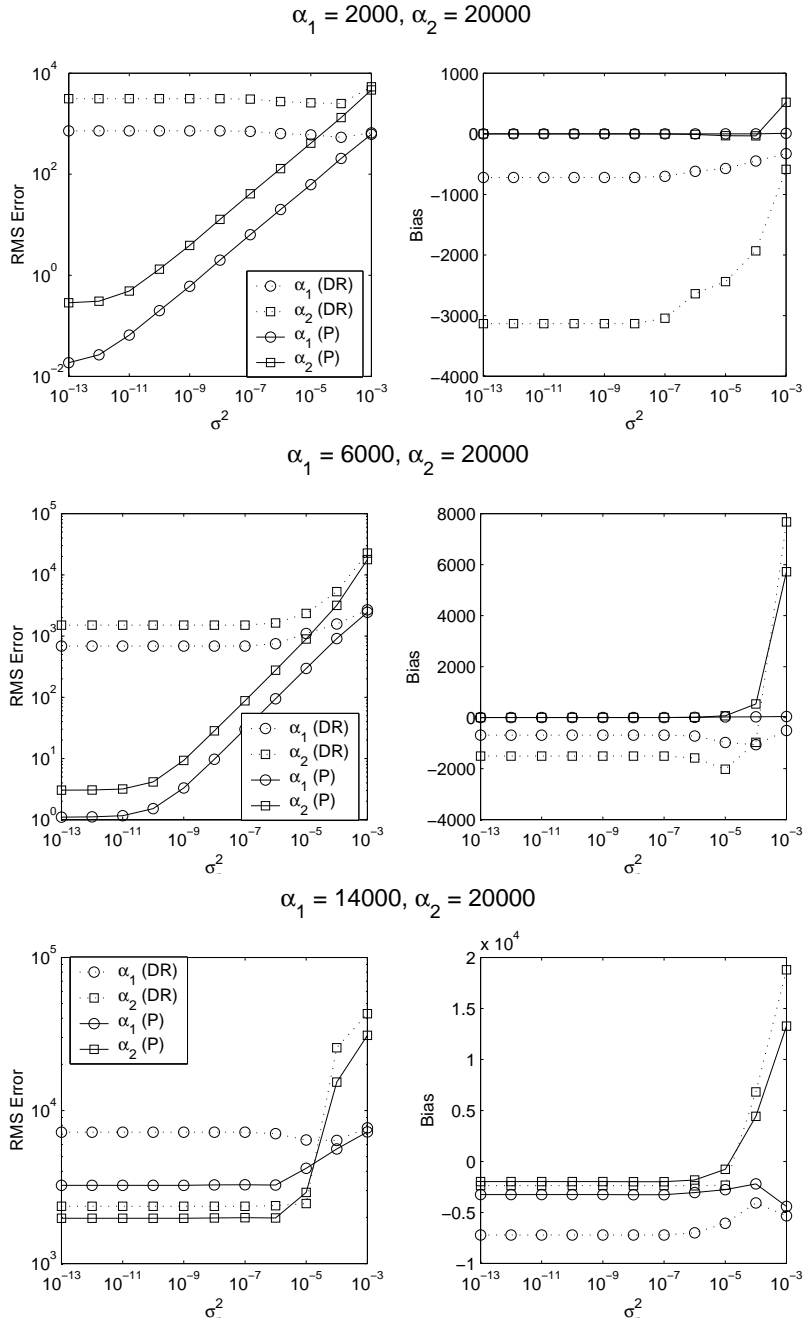


Figure 7. Simulation results for decay rates of $\{2000, 20000\}$, $\{6000, 20000\}$, $\{14000, 20000\}$. DR indicates decay rate parameterization of the signal model and P indicates pole parameterization of the signal model.




Figure 8. ROC performance of a correlation detector and the HMM on Yuma data collected with the Boom-SAR.




Figure 9. ROC performance for Bayesian algorithms during the JPG-IV demonstration using ad-hoc and physics-based feature sets.




Figure 10. ROC performance for MTADS JPG IV data using leave-one-out training.

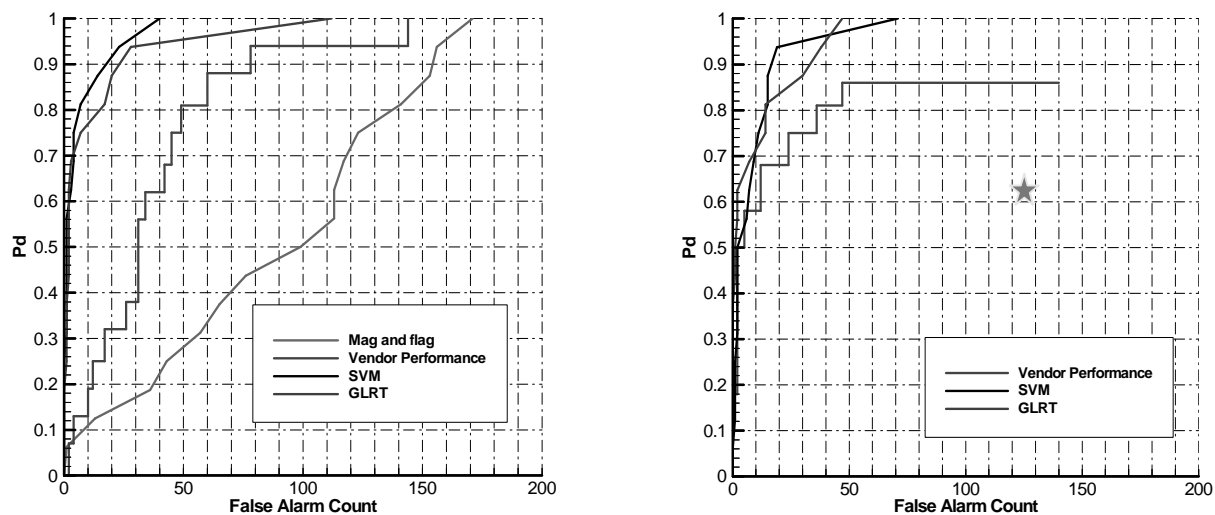


Figure 11. ROC Performance for JPG V, Area 3. Left panel, GEM-3 results on 202 anomalies, 16 UXO. Right panel, MTADS results on 141 anomalies, 16 UXO.

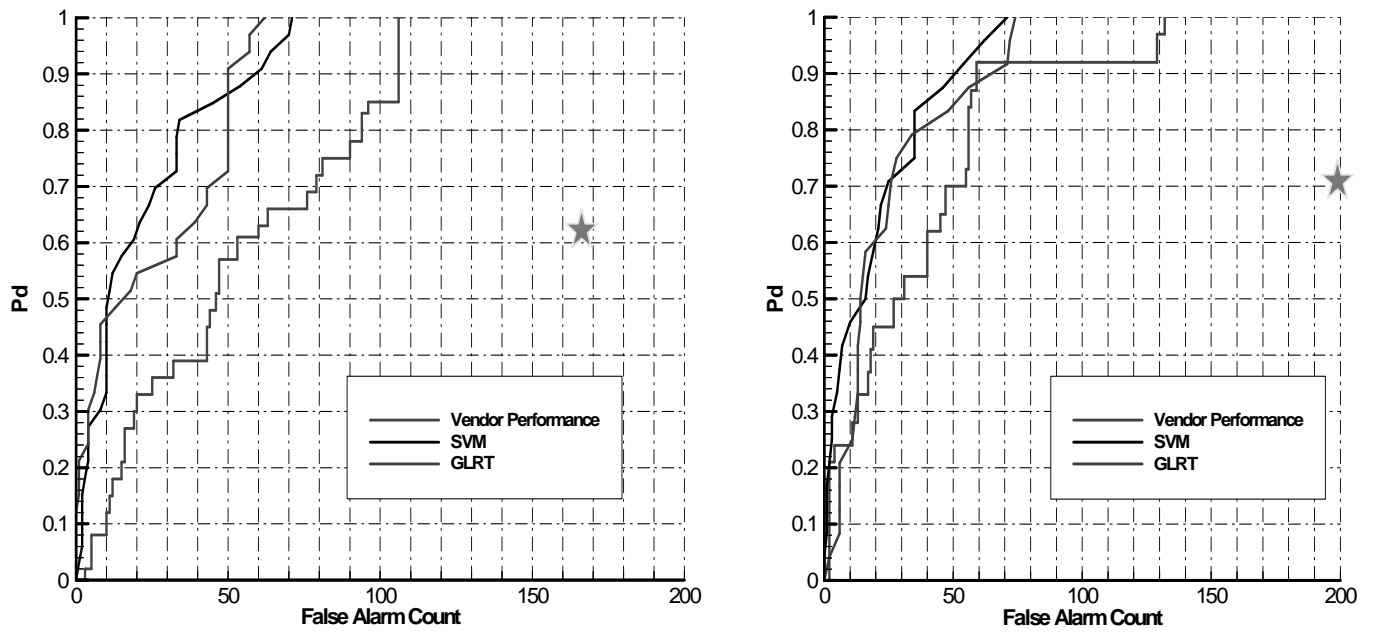


Figure 12. ROC Performance for JPG V, GEM3 data scored blind. Left panel, Area 1 results on 178 anomalies, 33 UXO. Right panel, Area 2 results on 236 anomalies, 24 UXO.





Figure 13. ROC Performance for JPG V, MTADS data scored blind. Left panel, Area 1 results on 149 anomalies, 33 UXO. Right panel, Area 2 results on 147 anomalies, 24 UXO.